

Supplementary Material

(not intended for publication)

All aboard: The effects of port development

César Ducruet Réka Juhász Dávid Krisztián Nagy Claudia Steinwender

Table of Contents

I	Theory	2
I.1	Equilibrium of the model	2
I.2	Equilibrium land use, wages, city populations and shipping flows . .	4
I.3	Inverting the model	10
I.4	Counterfactual simulation	11
I.5	Benchmark models used to decompose the aggregate welfare effects of containerization	12
I.6	A model with labor used in transshipment	15
I.7	A model with monopolistic competition in transshipment	19
II	Data	21
II.1	Lloyd's List shipping data	21
II.2	Underwater elevation levels	21
II.3	Saiz proxy for land-rents	22
II.4	Predicted city-level GDP per capita	22
II.5	Port shares for 1990	23
II.6	Area per throughput calculation for the Port of Seattle	25
II.7	Google Earth port area and containerization, modern data	25
II.8	Land reclamation	27
II.9	Country GDP per capita	27
II.10	Identifying city centroids for within port-city moves	28
II.11	Ship size data	28
II.12	Ship positions	29
II.13	Nautical maps for dredging dummy variable	29
II.14	Port cost data based on Blonigen and Wilson (2008)	30
II.15	Data on frost-free days	31
II.16	Annual reports for ports	31
II.17	Maritime Silkroad Counterfactual, regression with country fixed effects	32

I Theory

This section provides supplementary material to Sections 4.1, 6 and 8. Sections I.1 and I.2 supplement Section 4.1 by defining the equilibrium of the model as well as deriving the model equations that characterize cities' equilibrium land use, wages, populations and shipping flows, respectively. Section I.3 supplements Section 6.1 by showing how we invert the equilibrium conditions to back out amenities, productivities and exogenous port costs as a function of observed population, wages and the value of shipments. Section I.4 describes how we simulate the model for the counterfactuals of Sections 6 and 8. Section I.5 describes the benchmark models we use to decompose the aggregate welfare effects of containerization in Section 6.4. Finally, Sections I.6 and I.7 present two extensions to the baseline model of Section 4.1: one in which transshipment requires both labor and land (Section I.6), and one in which landlords engage in monopolistic competition in the transshipment sector (Section I.7).

I.1 Equilibrium of the model

We define the equilibrium of the model as follows.

Definition 1. *Given structural parameters $\alpha, \gamma, \eta, \sigma, \theta, \lambda$, the number of cities S and the subset of port cities $P \subseteq \{1, \dots, S\}$, country populations N_c , city amenities $a : \{1, \dots, S\} \rightarrow \mathbb{R}$, productivities $A : \{1, \dots, S\} \rightarrow \mathbb{R}$, exogenous transshipment costs $\nu : P \rightarrow \mathbb{R}$, inland and sea shipping costs as a function of distance $\phi_\zeta, \phi_\tau : \mathbb{R} \rightarrow \mathbb{R}$ and endogenous transshipment costs as a function of port share $\psi : (0, 1) \rightarrow \mathbb{R}$, an **equilibrium** of the model is a set of city populations $N : S \rightarrow \mathbb{R}$, nominal wages $w : S \rightarrow \mathbb{R}$, land rents $R : S \rightarrow \mathbb{R}$, employment levels $n : S \rightarrow \mathbb{R}$, port shares $F : S \rightarrow [0, 1)$, port-level shipping flows $Shipping : P \rightarrow \mathbb{R}$, the prices of transshipment services $O : P \rightarrow \mathbb{R}$, the prices of goods $p : S^2 \rightarrow \mathbb{R}$ and the quantities of goods $q : S^2 \rightarrow \mathbb{R}$ such that*

1. *workers choose their consumption of goods and city of residence within their country to maximize their utility (3), taking prices and wages as given;*
2. *landlords in each city r choose their consumption of goods and land use to maximize their utility*

$$u_L(r) = \left[\sum_{s=1}^S q_L(s, r)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \quad (\text{S.1})$$

*taking prices, land rents and shipping flows as given;*¹

¹We assume that landlords do not enjoy city amenities and do not have idiosyncratic tastes

3. *competition among landlords drives the price of transshipment services down to marginal cost, (4), and landlords' profits from transshipment down to zero;*²
4. *firms in each city r choose their production, employment and land use to maximize their profits*

$$\max_{n(r), 1-F(r)} p(r, r) \tilde{A}(r) n(r)^\gamma (1 - F(r))^{1-\gamma} - w(r) n(r) - R(r) (1 - F(r)) \quad (\text{S.2})$$

taking prices, land rents and wages as given, where $p(r, r)$ is the factory gate price of the good produced by the firm, and choose the shipping route to each destination to maximize their profits;

5. *competition among firms drives their profits down to zero;*
6. *there is no possibility of arbitrage, implying that the price of good r at s equals the expected iceberg cost over the factory gate price,*

$$p(r, s) = p(r, r) \mathbf{E}[T(r, s)]; \quad (\text{S.3})$$

7. *the market for labor clears in each city r , implying $n(r) = N(r)$;*
8. *national labor markets clear, implying $\sum_{r \in c} N(r) = N_c$ in each country c ;*
9. *the market for land clears in each city;*
10. *the market for transshipment services clears in each port city;*
11. *the market for each good clears worldwide.*

Note that this equilibrium definition implies that we do not give landlords the right to choose the amount of transshipment they conduct. In other words, landlords cannot refuse the provision of transshipment services to anyone at the market price. This assumption is needed for computational tractability, as it allows us to abstract from a corner solution in which the supply of transshipment services is zero. In line with this logic, we can relax the assumption and allow landlords to choose any *positive* amount of transshipment, but not zero transshipment. Generalizing the model this way does not change the equilibrium as landlords' profits are linear in the amount of transshipment for cities. As landlords are immobile, this assumption does not have any consequence on their optimal choices and is therefore without loss of generality.

²We relax this assumption in the monopolistic competition version of the model, presented in Section I.7.

and zero in equilibrium, hence landlords are indifferent between transshipping any two amounts as long as they are both positive.³

I.2 Equilibrium land use, wages, city populations and shipping flows

This section uses the equilibrium conditions of Section I.1 to characterize cities' equilibrium land use, wages, populations and shipping flows. To obtain these, we proceed as follows. Section I.2.1 solves for workers' optimal location choices. Section I.2.2 solves the landlords' problem for the optimal allocation of land between production and transshipment. Section I.2.3 solves the firms' problem, while Section I.2.4 uses equilibrium prices, the price index and market clearing to obtain the equations characterizing cities' equilibrium wages and population. Finally, Section I.2.5 derives the value of shipments flowing through any port in equilibrium.

I.2.1 Workers' optimal location choices

The utility function of workers, (3), implies that the indirect utility of a worker living in city r equals

$$u_j(r) = \frac{w(r)}{P(r)} a(r) b_j(r)$$

where $w(r)$ is the nominal wage and $P(r)$ is the CES price index of consumption goods in the city.

We assume that $b_j(r)$ is distributed Fréchet with scale parameter one and shape parameter $1/\eta$:

$$Pr(b_j(r) \leq b) = e^{-b^{-1/\eta}}$$

from which we obtain that the worker's indirect utility is also distributed Fréchet with scale parameter $\left[\frac{w(r)}{P(r)} a(r)\right]^{1/\eta}$:

$$Pr(u_j(r) \leq u) = e^{-\left[\frac{w(r)}{P(r)} a(r)\right]^{1/\eta} u^{-1/\eta}}$$

and hence, by the properties of the Fréchet distribution, the probability with which a worker chooses to live in city r is given by

$$Pr(u_j(r) \geq u_j(s) \ \forall s \neq r) = \frac{\left[\frac{w(r)}{P(r)} a(r)\right]^{1/\eta}}{\sum_{s \in c} \left[\frac{w(s)}{P(s)} a(s)\right]^{1/\eta}}.$$

³In the monopolistic competition version of the model (Section I.7), we do not need to make this assumption. In that model, landlords have market power and therefore choose both the price and the quantity of transshipment in a way that maximizes their profits.

In equilibrium, the fraction of workers choosing to live in city r coincides with this probability, implying

$$\frac{N(r)}{\sum_{s \in c} N(s)} = \frac{\left[\frac{w(r)}{P(r)} a(r) \right]^{1/\eta}}{\sum_{s \in c} \left[\frac{w(s)}{P(s)} a(s) \right]^{1/\eta}}. \quad (\text{S.4})$$

1.2.2 Landlords' optimal land use

Landlords earn income from providing transshipment services and from renting out land to firms that produce the city-specific good. Their utility function, (S.1), implies that the indirect utility of a landlord in city r equals her nominal income divided by the price index,

$$u_L(r) = \frac{\left[O(r) - (\nu(r) + \psi(F(r))) \text{Shipping}(r)^\lambda \right] \text{Shipping}(r) + R(r)(1 - F(r))}{P(r)}$$

where $O(r)$ is the price of transshipment services in city r (taken as given by the landlord), $\nu(r)$ is the exogenous part of transshipment costs, $F(r)$ is the share of land allocated to the port, $\text{Shipping}(r)$ is the value of shipments flowing through the port, excluding the price of transshipment services (hence, total demand for transshipment services, again taken as given by the landlord), $R(r)$ is the land rent prevailing in the city, and $1 - F(r)$ is the share of land rented out to firms. That is, the first term in the numerator corresponds to the landlord's net nominal income from providing transshipment services, while the second term corresponds to her nominal income from renting out land to firms.

The landlord decides on the allocation of land, captured by the single variable $F(r)$, to maximize her utility. As she cannot influence the price index $P(r)$, this is equivalent to maximizing her nominal income:

$$\max_{F(r)} \left[O(r) - (\nu(r) + \psi(F(r))) \text{Shipping}(r)^\lambda \right] \text{Shipping}(r) + R(r)(1 - F(r))$$

The first-order condition to this maximization problem is

$$-\psi'(F(r)) \text{Shipping}(r)^{1+\lambda} - R(r) = 0$$

from which, by rearranging,

$$-\psi'(F(r)) = \frac{R(r)}{\text{Shipping}(r)^{1+\lambda}}. \quad (\text{S.5})$$

1.2.3 Firms' problem

Recall that the representative firm operating in city r faces the production function

$$q(r) = \tilde{A}(r) n(r)^\gamma (1 - F(r))^{1-\gamma}$$

and maximizes its profits, (S.2), by choosing its employment and land use. The first-order conditions to the firm's profit-maximization problem imply

$$R(r) = \frac{1-\gamma}{\gamma} \frac{w(r) N(r)}{1 - F(r)} \quad (\text{S.6})$$

where we have used labor market clearing, which implies $n(r) = N(r)$. Plugging this back into the firm's cost function and production function, we obtain that the firm's marginal cost of production is equal to

$$\gamma^{-\gamma} (1-\gamma)^{-(1-\gamma)} \tilde{A}(r)^{-1} w(r)^\gamma R(r)^{1-\gamma}$$

which, by perfect competition among firms, equals the factory gate price in equilibrium:

$$p(r, r) = \gamma^{-1} A(r)^{-1} (1 - F(r))^{-(1-\gamma)} N(r)^{1-\gamma-\alpha} w(r) \quad (\text{S.7})$$

where we have used (S.6) again, together with the fact that $\tilde{A}(r) = A(r) N(r)^\alpha$.

Finally, equation (S.6) also implies that total factor payments in city r equal

$$Y(r) = w(r) N(r) + R(r) (1 - F(r)) = w(r) N(r) + \frac{1-\gamma}{\gamma} w(r) N(r) = \frac{1}{\gamma} w(r) N(r). \quad (\text{S.8})$$

1.2.4 Equilibrium wages and populations

From the workers' and landlords' problems, we can derive the constant-elasticity demand for the city- r good in city s as

$$q(r, s) = p(r, s)^{-\sigma} P(s)^{\sigma-1} Y(s)$$

where $p(r, s)$ is the price paid by the consumer, which includes the shipping cost between r and s . Demand in value terms is equal to

$$p(r, s) q(r, s) = p(r, r)^{1-\sigma} P(s)^{\sigma-1} Y(s) \mathbf{E}[T(r, s)]^{1-\sigma}$$

where we have used equation (S.3).

Market clearing for the good produced in city r implies that total factor payments

in r equal worldwide demand for the good (in value terms):

$$\frac{1}{\gamma} w(r) N(r) = \sum_{s=1}^S p(r, r)^{1-\sigma} P(s)^{\sigma-1} \frac{1}{\gamma} w(s) N(s) \mathbf{E}[T(r, s)]$$

where we have used equation (S.8) to substitute for total factor payments on both sides.

Plugging (S.7) into this equation yields

$$\begin{aligned} w(r) N(r) &= \gamma^{\sigma-1} A(r)^{\sigma-1} (1 - F(r))^{(1-\gamma)(\sigma-1)} N(r)^{-(1-\gamma-\alpha)(\sigma-1)} \\ &\quad w(r)^{1-\sigma} \sum_{s=1}^S P(s)^{\sigma-1} w(s) N(s) \mathbf{E}[T(r, s)]^{1-\sigma}. \end{aligned} \quad (\text{S.9})$$

The CES price index in city r takes the form

$$P(r)^{1-\sigma} = \sum_{s=1}^S p(s, r)^{1-\sigma} = \sum_{s=1}^S p(s, s)^{1-\sigma} \mathbf{E}[T(s, r)]^{1-\sigma}.$$

Plugging factory gate prices (S.7) into this equation yields

$$P(r)^{1-\sigma} = \gamma^{\sigma-1} \sum_{s=1}^S A(s)^{\sigma-1} (1 - F(s))^{(1-\gamma)(\sigma-1)} w(s)^{1-\sigma} N(s)^{-(1-\gamma-\alpha)(\sigma-1)} \mathbf{E}[T(s, r)]^{1-\sigma}. \quad (\text{S.10})$$

Rearranging equation (S.4) yields the following expression for the price index:

$$P(r) = \tilde{a}(r) w(r) N(r)^{-\eta} \quad (\text{S.11})$$

where $\tilde{a}(r)$ can be obtained by scaling amenities $a(r)$ according to

$$\tilde{a}(r) = \aleph_c a(r) = \left[\frac{\sum_{s \in c} N(s)}{\sum_{s \in c} \left[\frac{w(s)}{P(s)} a(s) \right]^{1/\eta}} \right]^\eta a(r).$$

Plugging equation (S.11) into (S.9) yields

$$\begin{aligned} A(r)^{1-\sigma} (1 - F(r))^{-(1-\gamma)(\sigma-1)} w(r)^\sigma N(r)^{1+(1-\gamma-\alpha)(\sigma-1)} &= \\ \gamma^{\sigma-1} \sum_{s=1}^S \tilde{a}(s)^{\sigma-1} w(s)^\sigma N(s)^{1-\eta(\sigma-1)} \mathbf{E}[T(r, s)]^{1-\sigma} \end{aligned} \quad (\text{S.12})$$

while plugging equation (S.11) into (S.10) yields

$$\begin{aligned} \tilde{a}(r)^{1-\sigma} w(r)^{1-\sigma} N(r)^{\eta(\sigma-1)} &= \gamma^{\sigma-1}. \\ \sum_{s=1}^S A(s)^{\sigma-1} (1-F(s))^{(1-\gamma)(\sigma-1)} w(s)^{1-\sigma} N(s)^{-(1-\gamma-\alpha)(\sigma-1)} \mathbf{E}[T(s, r)]^{1-\sigma} &. \end{aligned} \quad (\text{S.13})$$

Note that our assumptions on trade costs guarantee symmetry and hence $\mathbf{E}[T(r, s)]^{1-\sigma} = \mathbf{E}[T(s, r)]^{1-\sigma}$. Given this, we can show that equations (S.12) and (S.13) can be simplified further. To see that this is the case, guess that wages take the form

$$w(r) = \tilde{a}(r)^{\iota_1} A(r)^{\iota_2} (1-F(r))^{\iota_3} N(r)^{\iota_4}.$$

That is, they only depend on local amenities, productivity, land available for production, and population. Inspecting equations (S.12) and (S.13), one can verify that this guess is indeed correct if

$$\begin{aligned} \iota_1 &= -\frac{\sigma-1}{2\sigma-1}, \\ \iota_2 = \iota_3 &= (1-\gamma) \frac{\sigma-1}{2\sigma-1} \end{aligned}$$

and

$$\iota_4 = [\eta - (1-\gamma)(1-\alpha)(\sigma-1) - 1] \frac{1}{2\sigma-1}$$

as (S.12) and (S.13) reduce to the same equation if the guess is correct with these values of $\iota_1, \iota_2, \iota_3$ and ι_4 . Thus, wages in city r are given by

$$w(r) = \tilde{a}(r)^{-\frac{\sigma-1}{2\sigma-1}} A(r)^{\frac{\sigma-1}{2\sigma-1}} (1-F(r))^{(1-\gamma)\frac{\sigma-1}{2\sigma-1}} N(r)^{[\eta-(1-\gamma-\alpha)(\sigma-1)-1]\frac{1}{2\sigma-1}}.^4 \quad (\text{S.14})$$

Finally, plugging (S.14) back into either (S.12) or (S.13) gives us an equation that determines the distribution of population across cities:

$$N(r)^{[1+\eta\sigma+(1-\gamma-\alpha)(\sigma-1)]\frac{\sigma-1}{2\sigma-1}} = \gamma^{\sigma-1} \tilde{a}(r)^{\frac{\sigma(\sigma-1)}{2\sigma-1}} A(r)^{\frac{(\sigma-1)^2}{2\sigma-1}} (1-F(r))^{(1-\gamma)\frac{(\sigma-1)^2}{2\sigma-1}} MA(r) \quad (\text{S.15})$$

where

$$MA(r) = \sum_{s=1}^S \frac{\tilde{a}(s)^{\frac{(\sigma-1)^2}{2\sigma-1}} A(s)^{\frac{\sigma(\sigma-1)}{2\sigma-1}} (1-F(s))^{(1-\gamma)\frac{\sigma(\sigma-1)}{2\sigma-1}} N(s)^{[1-\eta(\sigma-1)-(1-\gamma-\alpha)\sigma]\frac{\sigma-1}{2\sigma-1}}}{\mathbf{E}[T(r, s)]^{\sigma-1}}$$

⁴We can freely choose the intercept of this equation as we have not normalized any price yet. We choose it to be equal to one.

is the market access of city r .

1.2.5 Equilibrium shipping flows

This section derives the equilibrium value of shipping flows through any port. To obtain these, we first need to introduce further notation. Let Z be an $S + P$ by $S + P$ matrix, where P denotes both the set and the number of ports in the model.⁵ Each of the first S rows and columns of Z corresponds to a city, while each of the last P rows and columns of Z corresponds to a port. Let us call a city or a port a *location*; that is, each row and column in Z corresponds to one location. We assume that an entry $z(i, \ell)$ of Z is zero if locations i and ℓ are not directly connected, or if $i = \ell$. Otherwise, $z(i, \ell)$ is defined as

$$z(i, \ell) = [\bar{T}(i, \ell) [1 + O(\ell)]]^{-\theta}$$

where $\bar{T}(i, \ell)$ is the common cost of shipping from i to ℓ directly, and $O(\ell)$ is the price of transshipment services at ℓ . If ℓ is a port belonging to port city r , then this price is given by equation (4). If ℓ is not a port but a (port or non-port) city, then we define $O(\ell) = 0$.⁶

Following Allen and Arkolakis (2019), we can show that the expected cost of shipping from city r to s can be written as

$$\mathbf{E}[T(r, s)] = \Gamma\left(\frac{\theta + 1}{\theta}\right) x(r, s)^{-1/\theta}$$

where $x(r, s)$ is the (r, s) entry of the matrix

$$X = (I - Z)^{-1}$$

and I is the $S + P$ by $S + P$ identity matrix.

Similarly, we can show that, if a good is shipped from city r to s , the probability that it is shipped through port k is given by

$$\pi(k|r, s) = \frac{x(r, k) x(k, s)}{x(r, s)}. \quad (\text{S.16})$$

and therefore the total value of goods shipped through port k from city r to city s

⁵Recall that S is the total number of (port or non-port) cities.

⁶For computational reasons, we need to add a small iceberg cost of shipping between each port and its own city. This cost equals 1.03 in both the inversion and the model simulations.

(excluding the price paid for transshipment services at k) equals

$$Shipping(k|r, s) = [1 + O(k)]^{-1} p(r, s)^{1-\sigma} P(s)^{\sigma-1} \frac{1}{\gamma} w(s) N(s) \pi(k|r, s).$$

Combining this with equations (S.3), (S.7), (S.11) and (S.16) yields

$$Shipping(k|r, s) = \gamma^{\sigma-2} [1 + O(k)]^{-1} A(r)^{\sigma-1} (1 - F(r))^{(1-\gamma)(\sigma-1)} N(r)^{-(1-\alpha-\gamma)(\sigma-1)}.$$

$$w(r)^{1-\sigma} \tilde{a}(s)^{\sigma-1} N(s)^{1-\eta(\sigma-1)} w(s)^\sigma \mathbf{E}[T(r, s)]^{1-\sigma} \frac{x(r, k) x(k, s)}{x(r, s)}$$

and therefore the total value of shipping through port k is given by

$$Shipping(k) = \gamma^{\sigma-2} [1 + O(k)]^{-1} \sum_r D_1(r) x(r, k) \sum_s D_2(s) \frac{\mathbf{E}[T(r, s)]^{1-\sigma}}{x(r, s)} x(k, s) \quad (\text{S.17})$$

where

$$D_1(r) = A(r)^{\sigma-1} (1 - F(r))^{(1-\gamma)(\sigma-1)} N(r)^{-(1-\alpha-\gamma)(\sigma-1)} w(r)^{1-\sigma}$$

and

$$D_2(s) = \tilde{a}(s)^{\sigma-1} N(s)^{1-\eta(\sigma-1)} w(s)^\sigma.$$

I.3 Inverting the model

This section describes how we invert the equilibrium conditions of the model to back out amenities, productivities and exogenous transshipment costs as a function of observed population, wages and the value of shipments. As a first step, we use the observed data to back out port shares in the model. To this end, we combine equations (S.5) and (S.6) to obtain port shares as a function of wages $w(r)$, population $N(r)$ and the value of shipments $Shipping(r)$ in each port city r :

$$-\psi'(F(r))(1 - F(r)) = \frac{1 - \gamma}{\gamma} \frac{w(r) N(r)}{Shipping(r)^{1+\lambda}} \quad (\text{S.18})$$

Given the assumptions we made on ψ' , the left-hand side of equation (S.18) is strictly decreasing in $F(r)$. Moreover, the left-hand side takes every real value between zero and infinity as ψ' is continuous, $\lim_{F \rightarrow 1} \psi'(F) = 0$ and $\lim_{F \rightarrow 0} \psi'(F) = -\infty$. This guarantees that solving equation (S.18) identifies a unique value of $F(r) \in (0, 1)$ for every port city.

The second step consists of solving for $\tilde{a}(r)$, $A(r)$ and $\nu(r)$ for the observed $N(r)$, $w(r)$ and $Shipping(r)$, as well as the $F(r)$ recovered in the previous step. This is done

using an algorithm that consists of an outer loop and an inner loop. In the inner loop, we obtain the values of $\tilde{a}(r)$ that solve the system of equations

$$\tilde{a}(r)^{1-\sigma} w(r)^{1-\sigma} N(r)^{\eta(\sigma-1)} = \gamma^{\sigma-1} \sum_{s=1}^S \tilde{a}(s)^{\sigma-1} w(s)^{\sigma} N(s)^{1-\eta(\sigma-1)} \mathbf{E}[T(r, s)]^{1-\sigma} \quad (\text{S.19})$$

derived from equations (S.12) and (S.13) for a *fixed* set of exogenous transshipment costs $\nu(r)$, and hence for fixed $\mathbf{E}[T(r, s)]$. For any $\mathbf{E}[T(r, s)]$, this system yields a unique solution for $\tilde{a}(r)$. Rearranging equation (S.14), we can then uniquely express productivity $A(r)$ as a function of the recovered $\tilde{a}(r)$:

$$A(r) = \tilde{a}(r) (1 - F(r))^{\gamma-1} w(r)^{\frac{2\sigma-1}{\sigma-1}} N(r)^{-[\eta-(1-\gamma-\alpha)(\sigma-1)-1]\frac{1}{\sigma-1}} \quad (\text{S.20})$$

In the outer loop, we search for the set of $\nu(r)$ for which the value of shipments implied by equation (S.17) – hence, by $N(r)$, $w(r)$, $F(r)$ and the recovered $\tilde{a}(r)$ and $A(r)$ – rationalize the shipping flows observed in the data. In practice, we start from a uniform guess of $\nu(r) = \bar{\nu}$, then perform a large number of iterations in which we update $\nu(r)$ gradually to get closer to satisfying equation (S.17). We also update $\mathbf{E}[T(r, s)]$ in every iteration step. Even though we cannot prove that this procedure identifies a unique set of $\nu(r)$, the algorithm has been converging to the same fixed point for various different initial guesses on $\nu(r)$, even when guessing non-uniform distributions of $\nu(r)$ initially.

I.4 Counterfactual simulation

This section describes how we perform counterfactual simulations in the model. First, we need to choose the absolute level of amenities $a(r)$ in each city r , as the inversion only identifies amenities up to a country-level scale, $\tilde{a}(r) = \aleph_c a(r)$. Unfortunately, nothing in the data guides us with this choice. Hence, we make the simplest possible assumption by assuming that average amenities are the same across countries and are equal to one:

$$\frac{1}{C_c} \sum_{r \in c} a(r) = \frac{1}{C_c} \sum_{r \in c} \frac{\tilde{a}(r)}{\aleph_c} = 1$$

where C_c denotes the number of cities in country c . Rearranging yields

$$\aleph_c = \frac{1}{C_c} \sum_{r \in c} \tilde{a}(r)$$

and hence we can obtain the absolute level of amenities in each city r as

$$a(r) = \frac{\tilde{a}(r)}{\aleph_c} = \frac{C_c}{\sum_{s \in c} \tilde{a}(s)} \tilde{a}(r).$$

Second, we solve for the counterfactual equilibrium of the model using an algorithm that consists of three loops embedded in each other. In the innermost loop, we obtain the distribution of population $N(r)$ that solves equation (S.15) for a *fixed* set of \aleph_c , $F(r)$ and $Shipping(r)$ (implying that $\mathbf{E}[T(r, s)]$ are also fixed). For any \aleph_c , $F(r)$ and $Shipping(r)$, equation (S.15) can be shown to have a unique positive solution if

$$\alpha < 1 - \gamma + \eta$$

which holds under the assumptions made in Section 4.1. Moreover, the solution can be obtained by simply iterating on equation (S.15), starting from any initial guess on $N(r)$. The proof of these results follows directly from the proof of equilibrium uniqueness in Allen and Arkolakis (2014).

In the middle loop, we solve for the set of country-specific \aleph_c that guarantee that the sum of city populations equals total country population in each country:

$$\sum_{r \in c} N(r) = N_c$$

where N_c denotes the exogenously given population of country c . We also solve for wages using equation (S.14) and for rents using equation (S.6).

In the outermost loop, we iterate on the distribution of port shares and shipping flows that satisfy both equations (S.5) and (S.17), also updating $\mathbf{E}[T(r, s)]$ in every step. We use the distributions of port share and shipping obtained in the inversion as our initial guesses. Even though we cannot prove that this procedure yields a unique equilibrium, we have been converging to the same distribution of endogenous variables for different initial guesses as well.

I.5 Benchmark models used to decompose the aggregate welfare effects of containerization

This section provides a description of the two benchmark models (Benchmark 1 and Benchmark 2) used to decompose the aggregate welfare gains from containerization.

1.5.1 Benchmark 1: No land use in transshipment

In Benchmark 1, we abstract from endogenous (land-dependent) transshipment costs. Thus, the cost of handling one unit of a good at port p_m is given by

$$\nu(p_m) \text{Shipping}(p_m)^\lambda$$

and, by perfect competition, the price of transshipment services equals this cost:

$$O(p_m) = \nu(p_m) \text{Shipping}(p_m)^\lambda \quad (\text{S.21})$$

As production is the only sector in which land can be productively used in this model, landlords optimally set the fraction of production land to one: $1 - F(r) = 1$. The remaining assumptions are the same as in the baseline model. Naturally, equation (S.5) does not hold in Benchmark 1, since all port shares are equal to zero.

Taking Benchmark 1 to the data. Taking Benchmark 1 to 1990 data follows similar steps as taking our baseline model to the data. We keep the structural parameters and the inland and sea shipping costs unchanged relative to the baseline model. To back out amenities, productivities and exogenous transshipment costs after containerization, we invert Benchmark 1 using 1990 data on population, wages and the value of shipments. This inversion procedure differs from the inversion of the baseline model in that we do not need to solve equation (S.5) for equilibrium port shares. As a result, we can skip the first step of the inversion procedure and immediately start with what we labeled as the second step in Section I.3.

In particular, we solve an algorithm that consists of an outer loop and an inner loop. In the inner loop, we obtain the values of city amenities $\tilde{a}(r)$ that solve equation (S.19), which holds in Benchmark 1 as well, for a *fixed* set of $\nu(r)$, hence for fixed $\mathbf{E}[T(r, s)]$. Once we have $\tilde{a}(r)$, we can obtain city productivities $A(r)$ from equation (S.20), which also holds in Benchmark 1, such that we set $1 - F(r) = 1$.

In the outer loop, we search for the set of $\nu(r)$ such that shipments implied by equation (S.17) equal the shipping flows observed in the data. Equation (S.17) also holds in Benchmark 1, except that we need to use $1 - F(r) = 1$ and equation (S.21) instead of equation (4) to calculate transshipment prices. In practice, we start from a uniform guess of $\nu(r) = \bar{\nu}$, then perform a large number of iterations in which we update $\nu(r)$ gradually to get closer to satisfying equation (S.17). We also update $\mathbf{E}[T(r, s)]$ in every iteration step.

Counterfactual simulation of Benchmark 1. When conducting the no-containerization counterfactual in Benchmark 1, we again try to stay as close as possible to our base-

line model. We offset the relationship between $\log \nu(r)$ and port depth, and increase all $\log \nu(r)$ by a constant ν_{CF} such that we have the same increase in international trade to world GDP (4.7 pp) as in the baseline model (Section 6.2). We also use the same procedure to obtain $a(r)$ from $\tilde{a}(r)$ (Section I.4). When conducting the targeted port development counterfactual, we decrease exogenous transshipment costs in the 24 targeted ports by 10%, as in the baseline model.

We solve for counterfactual equilibria using an algorithm that consists of three loops embedded in each other. In the innermost loop, we obtain the distribution of population $N(r)$ that solves equation (S.15) for a *fixed* set of \aleph_c and $Shipping(r)$ (implying that $\mathbf{E}[T(r, s)]$ are also fixed). Equation (S.15) is unchanged relative to the baseline model, except that we need to use $1 - F(r) = 1$. We follow the same iterative procedure as in Section I.4 to solve equation (S.15).

In the middle loop, we solve for the set of country-specific \aleph_c such that the sum of city populations equals total country population in each country. We also solve for wages using equation (S.14), which is the same as in the baseline model, except that $1 - F(r) = 1$.

In the outermost loop, we iterate on equation (S.17) to obtain equilibrium shipping flows, also updating $\mathbf{E}[T(r, s)]$ in every step. In contrast to the baseline model, we use $1 - F(r) = 1$ and equation (S.21) instead of equation (4) in this process. We use the 1990 shipping flows as our initial guess.

1.5.2 Benchmark 2: Land use in transshipment identical across port cities

In Benchmark 2, we allow for endogenous (land-dependent) transshipment costs. This implies that transshipment prices are given by equation (4), just like in our baseline model. However, we restrict transshipment land use to be identical across port cities. More precisely, we set the 1990 port share of each port city equal to the average 1990 port share in the baseline model. Similarly, we set the counterfactual port share equal to the average port share in the counterfactual of our baseline model. The remaining assumptions are the same as in the baseline model. Similar to Benchmark 1, equation (S.5) does not hold in this model since port shares are set exogenously through the above procedure, rather than optimally by port city landlords.

Taking Benchmark 2 to the data. We keep the structural parameters and the inland and sea shipping costs unchanged relative to the baseline model. To back out amenities, productivities and exogenous transshipment costs after containerization, we invert Benchmark 2 using 1990 data on population, wages and the value of shipments. Just like in Benchmark 1, we do not need to solve equation (S.5) for equilibrium port shares. As a result, we can skip the first step of the inversion procedure and immediately start from

the second step. This second step, in turn, is conducted exactly as in the baseline model (see Section I.3 for details), except that we use the average 1990 port share in the baseline model as $F(r)$ in each port city.

Counterfactual simulation of Benchmark 2. In the no-containerization counterfactual simulation of Benchmark 2, we change transshipment cost parameter β in the same way as in the counterfactual of the baseline model; offset the relationship between $\log \nu(r)$ and port depth; and increase all $\log \nu(r)$ by a constant ν_{CF} such that we have the same increase in international trade to world GDP (4.7 pp) as in the baseline model (Section 6.2). We also use the same procedure to obtain $a(r)$ from $\tilde{a}(r)$ (Section I.4).

Finally, we solve for the counterfactual equilibrium using an algorithm that consists of three loops embedded in each other. In the innermost loop, we obtain the distribution of population $N(r)$ that solves equation (S.15) for a *fixed* set of \aleph_c , $F(r)$ and $Shipping(r)$ (implying that $\mathbf{E}[T(r, s)]$ are also fixed). We use the average port share in the counterfactual of the baseline model as $F(r)$ in each port city. We follow the same iterative procedure as in Section I.4 to solve equation (S.15).

In the middle loop, we solve for the set of country-specific \aleph_c such that the sum of city populations equals total country population in each country. We also solve for wages using equation (S.14), which is the same as in the baseline model. We again use the same $F(r)$ in each port city.

In the outermost loop, we iterate on equation (S.17) to obtain equilibrium shipping flows, also updating $\mathbf{E}[T(r, s)]$ in every step. We again use the same $F(r)$ in each port city. We use the 1990 shipping flows as our initial guess.

I.6 A model with labor used in transshipment

This section presents a generalization of our baseline model in which the provision of transshipment services may require not only land, but also potentially labor. We show that, as long as the share of labor relative to land in transshipment is sufficiently low, this more general framework delivers predictions on port development and city populations that are extremely similar to the predictions of our baseline model. On the other hand, if the share of labor in transshipment is high, the model's predictions are in contrast with the empirical facts we document in Sections 3 and 5, as we describe below.

We now present the setup of the model with transshipment labor. Assume that the cost of transshipping one unit of a good in port city r equals

$$(\nu(r) + \psi(n^P(r)^{\gamma_P} F(r)^{1-\gamma_P})) Shipping(r)^\lambda$$

where $0 \leq \gamma_P \leq 1$. That is, γ_P is labor's share and $1 - \gamma_P$ is land's share in transship-

ment services. Our baseline model is a special case in which $\gamma_P = 0$. The remaining model assumptions are the same as in the baseline model.

We now show how our model predictions – more precisely, the three propositions of Section 4.2 – change in this more general framework. To obtain the first two propositions, note that the first-order conditions to the landlord's problem with respect to $n^P(r)$ and $F(r)$ together imply

$$n^P(r) = \frac{\gamma_P}{1 - \gamma_P} \frac{R(r)}{w(r)} F(r). \quad (\text{S.22})$$

On the production side, the first-order conditions to the firm's problem imply

$$n(r) = \frac{\gamma}{1 - \gamma} \frac{R(r)}{w(r)} (1 - F(r)). \quad (\text{S.23})$$

Adding equations (S.22) and (S.23) yields total demand for labor in the city,

$$N(r) = \frac{\gamma}{1 - \gamma} \frac{R(r)}{w(r)} (1 - \tilde{\gamma} F(r)) \quad (\text{S.24})$$

where $\tilde{\gamma} = \frac{\gamma/(1-\gamma) - \gamma_P/(1-\gamma_P)}{\gamma/(1-\gamma)}$. Combining equation (S.24) with equation (S.22), we obtain labor used for transshipment as

$$n^P(r) = \frac{\gamma_P}{1 - \gamma_P} \frac{1 - \gamma}{\gamma} N(r) \frac{F(r)}{1 - \tilde{\gamma} F(r)}$$

and hence the landlord's first-order conditions imply

$$-\psi' \left(\left[\frac{\gamma_P}{1 - \gamma_P} \frac{1 - \gamma}{\gamma} N(r) \right]^{\gamma_P} \frac{F(r)}{(1 - \tilde{\gamma} F(r))^{\gamma_P}} \right) = \hat{\gamma} \frac{w(r)^{\gamma_P} R(r)^{1-\gamma_P}}{\text{Shipping}(r)^{1+\lambda}} \quad (\text{S.25})$$

where $\hat{\gamma}$ is a constant. Equation (S.25) allows us to state the following two propositions.

Proposition 4. *Assume $\gamma_P \leq \gamma$. Then land allocated to the port is increasing in the amount of shipping flows.*

Proof. $\gamma_P \leq \gamma$ implies $\tilde{\gamma} > 0$. As a consequence, the argument inside the function $-\psi'$ is increasing in land allocated to the port, $F(r)$. Given the convexity of ψ , this means that the left-hand side of equation (S.25) is decreasing in $F(r)$. This, together with the fact that the right-hand side of (S.25) is decreasing in shipping flows $\text{Shipping}(r)$, yields the result. \square

Proposition 5. *Assume $\gamma_P \leq \gamma$. Then land allocated to the port is decreasing in land rents.*

Proof. The proof follows the exact same steps as the proof of Proposition 4. \square

Propositions 4 and 5 are the counterparts of Propositions 1 and 2 of Section 4.2. As the comparison of Propositions 4 and 5 to Propositions 1 and 2 clarifies, the sufficient condition under which the model with transshipment labor yields the same predictions as our baseline model is $\gamma_P \leq \gamma$. That is, labor's share in transshipment may be positive but needs to be below labor's share in the production of the city-specific good. This result is intuitive. Higher demand for transshipment, or a lower opportunity cost of transshipment, triggers an expansion of transshipment services in the city. As long as land's share in transshipment is higher than land's share in the rest of the economy, standard Heckscher–Ohlin logic dictates that this expansion is reached through more land used for transshipment and less in the rest of the economy.

If labor's share in transshipment is higher than labor's share in the production of the city-specific good, the model no longer yields clear-cut predictions on land allocation between the two sectors of the economy. In the extreme case in which land is not used in transshipment at all ($\gamma_P = 1$), port activity naturally does not depend on land rents whatsoever. This is clearly in contrast with our empirical facts documented in Section 3, and in particular, with the result that containerization increased shipping more in low land-rent cities.

To derive the counterpart of Proposition 3, note that land rents in the model with transshipment labor can be obtained from equation (S.24) as

$$R(r) = \frac{1 - \gamma}{\gamma} \frac{w(r) N(r)}{1 - \tilde{\gamma} F(r)}$$

whereas total income in city r is given by

$$\frac{1}{\gamma} w(r) n(r) = \frac{1}{\gamma} \frac{1 - F(r)}{1 - \tilde{\gamma} F(r)} w(r) N(r).$$

Using these results in the derivation of the equilibrium conditions, we obtain that the population of city r is the solution to the following equation:

$$N(r)^{[1+\eta\sigma+(1-\gamma-\alpha)(\sigma-1)]\frac{\sigma-1}{2\sigma-1}} = \gamma^{\sigma-1} \tilde{a}(r)^{\frac{\sigma(\sigma-1)}{2\sigma-1}} A(r)^{\frac{(\sigma-1)^2}{2\sigma-1}} (1 - \tilde{\gamma} F(r))^{[1+(1-\gamma)(\sigma-1)]\frac{\sigma-1}{2\sigma-1}} \cdot (1 - F(r))^{-\frac{\sigma-1}{2\sigma-1}} M A(r) \quad (\text{S.26})$$

where

$$MA(r) = \sum_{s=1}^S \tilde{a}(s)^{\frac{(\sigma-1)^2}{2\sigma-1}} A(s)^{\frac{\sigma(\sigma-1)}{2\sigma-1}} (1-F(s))^{\frac{\sigma-1}{2\sigma-1}} (1-\tilde{\gamma}F(s))^{[(1-\gamma)\sigma-1]\frac{\sigma-1}{2\sigma-1}} \cdot \\ N(s)^{[1-\eta(\sigma-1)-(1-\gamma-\alpha)\sigma]\frac{\sigma-1}{2\sigma-1}} \mathbf{E}[T(r,s)]^{1-\sigma}.$$

Equation (S.26) allows us to state the following proposition, which is the counterpart of Proposition 3 in Section 4.2.

Proposition 6. *If $\gamma_P < 1$, then an increase in the share of land allocated to the port in city in r , $F(r)$, decreases shipping costs $\mathbf{E}[T(r,s)]$, thus increasing $MA(r)$. Everything else fixed, an increase in $MA(r)$ increases the population of the city (market access effect). Holding $MA(r)$ fixed, if $\gamma_P \geq \gamma$, an increase in $F(r)$ draws additional people into the city (crowding-in effect). If $0 < \gamma_P < \gamma$, an increase in $F(r)$ may trigger either a crowding-in effect or migration out of the city (crowding-out effect), depending on the values of structural parameters γ , γ_P and σ . If and only if $\gamma_P = 0$ (our baseline model), the model implies a crowding-out effect irrespectively of the values of structural parameters.*

Proof. The results follow directly from equation (S.26). □

According to Proposition 6, an expansion of port activity has different implications on city population depending on labor's share in transshipment. Besides the standard market access effect, port development affects city population in two ways. First, it draws people into the transshipment sector as long as labor's share in the sector is different from zero. Second, it decreases the amount of land available for the production of the city-specific good, which induces workers in this sector to leave the city. If labor's share in the transshipment sector is sufficiently high, the first effect always dominates the second one (crowding in). This implies that the population of the city should increase even more than what is implied by the standard market access effect. Such a crowding-in effect, however, is not consistent what we find in the data (Section 5), in particular, with the negative and significant coefficient on shipping once we control for market access.

To sum up, the model presented in this section sheds light on two facts. First, if the share of labor in transshipment is too high, the model with transshipment labor has different implications than our baseline framework. These implications, however, are in clear contrast with the empirical findings of Sections 3 and 5. Second, if the share of labor in transshipment is sufficiently low, the model with transshipment labor is more complex in its structure but delivers predictions that are extremely similar to

the predictions of our baseline framework.

I.7 A model with monopolistic competition in transshipment

This section presents a version of our baseline model in which landlords providing transshipment services engage in monopolistic competition. This implies that, unlike in our baseline model, port activity involves positive profits. We also show how we take the model with monopolistic competition to the data and how we simulate the same no-containerization counterfactual in it as in our baseline model.

We first present the setup of the monopolistic competition model. As in our baseline model, we assume that each city is inhabited by a continuum of landlords. Without loss of generality, we normalize the mass of these landlords to one in each city, and index an individual landlord by $m \in [0, 1]$.

Unlike in our baseline model, we assume that transshipment services are differentiated products. Firms shipping through port city r may use the services of any number of landlords m residing in the city. Firms aggregate transshipment services in a CES function with elasticity of substitution $\zeta \in (1, \infty)$ across the services performed for them by the individual landlords. As $\zeta < \infty$, these services are imperfect substitutes. Hence, each firm uses the transshipment service of each landlord in equilibrium.⁷

Landlords are aware that they are the sole provider of their differentiated transshipment service but cannot influence city-wide prices and quantities. Thus, they engage in monopolistic competition, choosing their land allocation, transshipment price and transshipment quantity to maximize their net nominal income. In other words, landlord m in port city r solves the problem

$$\max_{F_m(r), O_m(r), Shipping_m(r)} \left[O_m(r) - (\nu(r) + \psi(F_m(r))) Shipping(r)^\lambda \right] Shipping_m(r) + R(r)(1 - F_m(r))$$

where $O_m(r)$ is the price of transshipment services that landlord m charges, $\nu(r)$ is the exogenous part of transshipment costs, $F_m(r)$ is the share of land that the landlord allocates to transshipment, $Shipping(r)$ is the total value of shipments flowing through the port excluding the price of transshipment services, $R(r)$ is the land rent prevailing in the city, and $1 - F_m(r)$ is the share of land rented out to firms.

As the price elasticity of demand for each landlord's transshipment service is constant at $-\zeta$, each landlord charges a constant markup over her marginal cost in equilib-

⁷To fix ideas, one may think that one port city landlord provides the cranes, another the storage, and so on. As a result, firms use the services of all landlords, not only one.

rium:

$$O_m(r) = \frac{\zeta}{\zeta - 1} (\nu(r) + \psi(F_m(r))) \text{Shipping}(r)^\lambda$$

As landlords in a given port city are symmetric, we can drop their index and simply write

$$O(r) = \frac{\zeta}{\zeta - 1} (\nu(r) + \psi(F(r))) \text{Shipping}(r)^\lambda \quad (\text{S.27})$$

from which we get that landlords earn profits on transshipment equal to

$$\Pi(r) = \frac{1}{\zeta - 1} (\nu(r) + \psi(F(r))) \text{Shipping}(r)^{1+\lambda}. \quad (\text{S.28})$$

For simplicity, we assume that landlords spend these profits outside our set of cities S . This implies that we do not need to take profits into account when calculating demand for goods in the city, or city GDP. This assumption helps us keep the model computationally tractable.

The first-order condition to the landlord's maximization problem with respect to $F_m(r)$ implies

$$-\psi'(F(r)) \text{Shipping}(r)^{1+\lambda} - R(r) = 0$$

from which, by rearranging,

$$-\psi'(F(r)) = \frac{R(r)}{\text{Shipping}(r)^{1+\lambda}}.$$

Note that this equation is identical to equation (5) of our baseline model. More generally, as the remaining model assumptions in the monopolistic competition model are the same as those in the baseline model, the only equation that differs between the two frameworks is equation (S.27), which replaces equation (4) in the baseline model. The remaining equilibrium conditions are all identical.

In Section A.3, we conduct a robustness check in which we take the model with monopolistic competition to the data to measure the aggregate gains from containerization, as in the baseline model. Inverting and simulating the monopolistic competition model follows the same steps as described in Sections I.3 and I.4, with one exception: we use equation (S.27) instead of equation (4) whenever we calculate transshipment prices.

To do so, we need to choose the value of the markup parameter ζ . Note that, by equation (S.28), transshipment profits are decreasing in ζ . Data on profits of ports are hard to find, especially during our period of interest, but we were able to obtain profit and revenue data for a number of ports from annual reports of port authorities between

1950 and 1990.⁸ In this sample, profits as a percentage of revenue are on average 28%, with no clear trends over time. Choosing $\zeta = 3$, our model predicts an average profit margin of 27% and a median profit margin of 33% across ports. Hence, we use $\zeta = 3$ in the inversion and the counterfactual simulation.⁹

II Data

In this section, we provide additional details about data construction and sources for the variables used in the analysis.

II.1 Lloyd’s List shipping data

We clean the shipping data by manually matching them to the 1953 and 2017 editions of the *World Port Index (WPI)*, which is a widely used reference list of worldwide ports. The initial Lloyd’s List sample of ‘ports’ included ports on navigable rivers such as Budapest, Hungary. We therefore chose to discipline the sample of ports using WPI. We use a historic and current edition of the WPI to ensure we capture both ports that may no longer exist, and ones that only appear later in the period. A different approach would have been to choose a distance threshold from the coast and drop any port located further from the coast than the threshold. This definition, however, is very sensitive to the precision of the coastline shapefile used to calculate distance from the coast, which is why we did not choose this method. Despite filtering the Lloyd’s List sample through the WPI, our final sample still contains a handful of ports that are very far inland. In the empirical analysis, we show that our results are robust to different ways of treating these ‘inland ports.’ Our base sample consists of Lloyd’s List ports that match to at least one of the WPI editions.

II.2 Underwater elevation levels

We use data on underwater elevation levels from the *General Bathymetric Chart of the Oceans (GEBCO)*. We use the 2014 version of these data. Most observations in the dataset are from ship-track soundings with interpolation between soundings guided by satellite-derived gravity data. The data are continuously updated with sources from local bathymetry offices and coastal navigation charts. More details on dataset construction can be found at <http://www.gebco.net>.

⁸We describe these data in Section II.16.

⁹We compute the profit margin of port r in the model as $\frac{\Pi(r) - R(r)F(r)}{O(r)Shipping(r)}$. These margins vary across ports and are in fact negative for a few of them. As these ports operate in the data, we do not let them shut down in the model and assume they are subsidized from the outside economy.

II.3 Saiz proxy for land-rents

The following sources are used to calculate the Saiz measure for our sample of cities. The coastline shapefile needed to distinguish between land and sea cells is from GSHHG (<https://www.soest.hawaii.edu/pwessel/gshhg/>). Inland bodies of water and wetlands are from the World Wildlife Fund’s *Global Lakes and Wetlands Database* (<https://www.worldwildlife.org/pages/global-lakes-and-wetlands-database>). Finally, data on land elevations used to calculate the slope of each cell is from GEBCO’s land data, described above.

II.4 Predicted city-level GDP per capita

Here we provide a more detailed discussion of how we estimate city level GDP per capita for our full sample of cities (port and non-port cities). First, we merge the *Canback* data with our city list, and construct GDP per capita from the level of GDP and the population data provided by Canback. GDP are reported at purchasing power parity (in 2005 USD). We have estimates from this source for 898 cities in our sample.

We estimate city-level GDP for the full sample by extrapolating the estimated relationship between GDP per capita and nightlight luminosity. We begin by estimating the linear fit of GDP per capita on nightlight luminosity, building on a growing body of evidence suggesting that income can be reasonably approximated using nightlight luminosity data (Donaldson and Storeygard, 2016).

We construct the ‘luminosity’ of each city in the following way. We take the 1992 30 arc-second grid layer from NOAA’s *National Geophysical Data Center* (source: <https://ngdc.noaa.gov/products/>) as the baseline input, as this is the closest year to 1990 – the year for which we have city income from *Canback*. We define a cell in this raster to be ‘lit up’ if its luminosity level is above 25. This threshold defines meaningful levels of economic activity in the cell - as proxied by nightlights.¹⁰ We then construct a polygon from contiguous cells with luminosity above 25 for each city in our sample. We observe luminosity for 2,294 cities in our dataset.¹¹ With these data in hand, we then define a city i ’s luminosity, $luminosity_i$, to be the sum of all cells’ luminosity levels within the polygon. Note that in this summation, we drop any cells identified as ‘gas flares’ in the source data, as these do not contain meaningful information on economic activity.

For the remaining 342 cities (13%), we either fail to identify an area polygon assigned to the city (340 cities) or a gas flare completely covers the polygon of the city (2 cities). We observe both GDP per capita and luminosity for a subset of 810 cities. For

¹⁰We experimented with different cutoffs and this was the one for which the R^2 in the regression of income on luminosity was highest.

¹¹We have cities with ‘missing’ luminosity data if we fail to detect *any* cells with luminosity levels above 25 in the vicinity of the city’s geocode.

this subset, we estimate the relationship between GDP per capita and luminosity. More precisely, we estimate

$$\ln(GDP/capita)_i = \beta * \ln(luminosity_i) + FE_c + \epsilon_i \quad (S.29)$$

where $GDP/capita_i$ is city-level GDP per capita as compiled in the *Canback Global Income Distribution Database (CGIDD)* for the year 1990 which covers 898 cities, and $luminosity_i$ measures the sum of luminosity in the cells in the polygon that defines the area of the city.

Note that most of the papers in this literature estimate the level of GDP within a country, where the level of development is not as widely dispersed as across cities worldwide. To account for these differences and the way in which they affect luminosity, we include country fixed effects FE_c in our estimation. However, in order to identify country fixed effects we need to drop 21 cities that are the only cities with GDP per capita data in their respective country, leaving a sample of 789 cities for estimation.

The results of this regression are given in column (1) of Supplementary Table S.1. We then predict GDP per capita for all cities for which we observe luminosity that are also in the set of countries used in this regression. This allows us to predict GDP per capita for a total of 2,289 cities. For the remaining 341 cities, we use the following approximation. For 89 cities, we observe GDP per capita directly, which we use. For 240 cities we only observe population in 1990, so we use this to predict GDP per capita based on the estimated relationship between GDP per capita and population in 1990 for all cities in our sample for which we observe both measures. This estimated relationship is given in column (2) of Supplementary Table S.1. Finally, for 18 cities we only observe population in 1980, so we use the latter to predict GDP per capita for all cities in our sample for which we observe both variables, resulting in the estimated relationship in column (3) of Supplementary Table S.1.

This procedure yields a city-level estimate for GDP per capita for all 2,636 cities in our dataset.

II.5 Port shares for 1990

Here, we provide details on the construction of port share data and the sources used. First, it is important to note that historical data on the area occupied by the port is very difficult to find. For example, data on port area is only sporadically and inconsistently reported in *Lloyd's Ports of the World*, and it is usually not found in ports' annual reports. These are in fact the two sources from which we take the measure for the ports where port area is observed. We also experimented with using satellite images from the 1980s, but the resolution is too low to detect port areas.

Table S.1: Relationship between GDP per capita and nightlight luminosity

Independent variables	ln(GDP per capita)		
	(1)	(2)	(3)
ln(Luminosity)	0.126*** (0.014)		
ln(Population, 1990)		0.107*** (0.013)	
ln(Population, 1980)			0.100*** (0.014)
Observations	789	854	871
R-squared	0.926	0.923	0.921
Country FE	✓	✓	✓

Notes: All regressions include country fixed effects. Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

We observe data on port area in 1990 for seven cities. These are: Aarhus (Denmark), Helsinki (Finland), Copenhagen (Denmark), Hamburg (Germany), Los Angeles (USA), New Orleans (USA) and Seattle (USA).¹² Data for the European ports and for the port of Los Angeles are from *Lloyd's Ports of the World (1990)*. We complemented these with data for other U.S. ports where planning maps and annual reports gave information on the land area of the port. In all these cases, we verified or cleaned the data to ensure that a consistent definition of port area was used. In particular, these measures only include the total land (and not sea) area occupied by the port. Data for the remaining U.S. ports are from *Port Authority of Seattle (1989)* and *Port of New Orleans (1984)*. These documents were shared by the port authorities based on requests we made. For Long Beach, we take port area in 1971 from the port's annual report (*Port of Long Beach, 1971*) and add additional land acquired from a detailed history of port projects (*Riffenburgh, 2012*). To construct the port shares, we use the area of *land* occupied by the city as reported in Wikipedia.

¹²The port area for Los Angeles includes the area occupied by the ports of Los Angeles and Long Beach.

II.6 Area per throughput calculation for the Port of Seattle

We obtained ‘Property Books’ that allow us to calculate the area of the Port of Seattle from the *Port of Seattle Public Records Office*. These volumes contain engineering maps for each parcel of land under the ownership of the port. Each map includes an estimate of the land area. For both years 1961 and 1973, we used only land parcels directly related to port activities. In particular, we excluded the airport and the marina terminal. Data on annual total throughput (in short tons, including both domestic and international sea-borne trade) and the share of containerized cargo were collected from *Annual Reports* that are archived at the *Puget Sound Regional Archives*. To smooth out fluctuations in year-to-year capacity utilization, we took the five-year moving-average of throughput.

Table S.2: Port of Seattle: area per unit of cargo shipped


Year	Area	Throughput	Area/Throughput
1961	8,651,016	2,022,192	4.28
1973	33,547,908	4,135,795	8.11

Notes: Area reported in square feet, throughput in short tons. Data were not available far enough back in time to allow for the calculation of the five-year moving-average for 1961.

Supplementary Table S.2 reports the numbers. While the expansion of traffic during this period was impressive (throughput doubled), the area occupied by the port expanded even more rapidly (increasing almost fourfold), such that area per throughput increased by 90% during this period. The *Annual Reports* paint a consistent picture. In the early 1960s, the port acquired vast parcels of land in the Lower Duwamish Industrial Development District. Throughout the latter half of the decade, the port continued to acquire more land in this area and to simultaneously develop the acquired tracts. These were completed in the late 1960s, early 1970s. We illustrate this in Supplementary Figure S.1 which shows the set of acquired land parcels and an example of a completed container facility.

II.7 Google Earth port area and containerization, modern data

We compiled data on the area of all ports for a random subset of port cities in our dataset (236 cities, which is 43% of the full sample), resulting in 252 individual ports. For each port, we hand-coded polygons that contain port activities based on satellite images from *Google Earth*. We used the name tags of buildings as well as visual markers (e.g., stacked containers, ships). We aimed to be conservative in that we only



Acquired land parcels (red shading), 1963



Completed terminal 102, 1970

Figure S.1: Illustration of port development, Seattle

Notes: The two panels illustrate development of the port through the 1960s. The first panel shows the initial set of land parcels acquired by the port along the Duwamish Waterway in the early 1960s. The second shows a container terminal completed in 1970 within this project. Sources: 'Port of Seattle: Industrial Development, Duwamish Waterway' (1963), 'Annual Report of the Port of Seattle' (1970).

included areas that could clearly be identified as containing port-related activities. As such, we did not include warehouses (as they cannot be unambiguously identified) or highways or railways. A port can have multiple polygons, e.g., in the case of terminals that are not directly connected. *Google Earth* reports the area (in km²) of each polygon, which we aggregate to the level of geopolis port cities. The average area of a port in our data is 3.6 km² (median: 2 km²), with a minimum of 0.03 km² and a maximum

of 30 km² (Los Angeles, including the Port of Long Beach). The latter occupies 43 km² according to Wikipedia (https://en.wikipedia.org/wiki/Port_of_Long_Beach and https://en.wikipedia.org/wiki/Port_of_Los_Angeles), so while our measure most likely underestimates the true size of ports, the measure is arguably in the correct range.

Data on total (in tons) and containerized (in TEUs; twenty-foot equivalent units) volume of cargo handled by each port is taken from the 2009 edition of *Le Journal de la Marine Marchande (JMM)*. We use the average of the reported numbers for 2008 and 2009 in order to maximize the number of observations, as some ports only report data for one of the two years. In order to generate the share of container traffic in total merchandise traffic, we use the average weight per TEU of 12 tons as recommended by the *European Sustainable Shipping Forum*.¹³

We match the dataset on the area of ports and cargo volume based on the names, countries and geocodes of the ports, resulting in 123 observations.

II.8 Land reclamation

Data on land reclaimed from the sea are taken from [Martín-Antón et al. \(2016\)](#). The authors compare historical maps to current Google Earth images to examine whether land reclamation has taken place in a city. We matched these data to our sample of port cities. The authors report three measures; i) any land reclamation, ii) coastal land reclamation, iii) coastal and island land reclamation. This contains land reclaimed for any purpose, not just for port activities. In our analysis, we use their coastal land reclamation measure, though the results are essentially the same regardless of the measure used.

The authors systematically examined the coastlines of the world, paying particular attention to South East Asia, the Persian Gulf, Europe and the U.S., where land reclamation has been more extensive. Any systematic measurement error introduced in this way will be accounted for in our specifications that control for continent and coastline fixed effects. Reassuringly, the coefficients of interest do not change substantially with the inclusion of these, suggesting that these issues – if present, are not quantitatively large.

II.9 Country GDP per capita

Data on country-level GDP per capita are from the *Penn World Tables*. We take real GDP at constant 2011 prices (USD) and divide by country population reported from the same source. In theory, the data exist for 1950 (our first sample year), but in practice there are many missing observations. For this reason, in robustness checks, we always

¹³Downloaded on March 11, 2021, from https://ec.europa.eu/clima/sites/clima/files/docs/0108/20170517_guidance_cargo_en.pdf.

use the data for 1960. This is observed for many, though not all countries.

II.10 Identifying city centroids for within port-city moves

In Section 3, we discussed evidence that showed that ports had moved further towards the outskirts of the city during our sample period. To conduct this exercise, we use data on ports' geocodes from two editions of the *World Port Index*: 1953 and 2017. We also need to identify the geocode of each city's centroid. To this end, we use daylight satellite data to identify a city's contiguous built-up area and find the city centroid within this polygon. We closely follow the methodology in Baragwanath, Goldblatt, Hanson, and Khandelwal (2019). In particular, we use an extremely high resolution dataset of daylight satellite data, the *Global Human Settlement Built-Up Grid* available at 38 m resolution (source: https://ghsl.jrc.ec.europa.eu/ghs_bu.php). Using this raster and the geocodes of our cities, we construct a polygon for each city consisting of contiguous built-up cells around the geocode. We take the centroid of this polygon to be the centroid of the city.

II.11 Ship size data

The evolution of ship sizes, illustrated in Supplementary Figure S.2 is based on data purchased from the *Miramar Ship Index* (Haworth, 2020), accessible at <http://www.miramarshipindex.nz>. The *Miramar Ship Index* is a comprehensive list of all newly built ships and their main characteristics going back to the 19th century. We calculate the average tonnage of all newly built ships in the years 1960, 1990, and 2010, distinguishing between container-ships and non-container ships.




Figure S.2: Development of ship sizes over time, 1960-2010

Notes: The figure illustrates the growth in ship size, as measured in average tons per newly built ship in a given year, for the years 1960, 1990, and 2010, for container-ships and all other ships (i.e., excluding container-ships), respectively. Source: Haworth (2020).

II.12 Ship positions

We purchased data on the precise geo-location of ships for 100 randomly selected ports in our sample from *marinetraffic.com*. Data were available for 94 of these 100 ports. The data refer to all stationary (i.e., reporting speed of 0 knots per hour) cargo vessels during the week of November 4 and 10, 2019, at 12:00-13:00 local time, resulting in 17,000 observations. For vessels that report different stationary positions during this one-hour window, we keep the last reported stationary location within the hour. We calculate the geodesic distance of each anchored vessel to the geocode of the port and take the sum across the number of anchored ships within certain distances from the geocode of the port. Supplementary Table S.3 shows the distribution of ships around the port by decile of port size.

Table S.3: Location of stationary ships around the port

Port	up to 1km	3km	5km	10km	15km	20km	25km	30km	Total
All	11	39	57	76	83	86	89	90	100
1st decile	27	91	100	100	100	100	100	100	100
2nd - 3rd decile	49	81	95	95	95	95	95	95	100
4th - 5th decile	26	64	88	99	100	100	100	100	100
6th - 7th decile	21	57	78	96	100	100	100	100	100
8th - 10th decile	8	33	50	70	79	83	86	88	100

Notes: The table shows the location of stationary cargo ships for 94 random ports in our sample. Data were requested for 100 random ports in our sample. Four ports we requested data for had no anchored ships during the time window when data were reported and two ports from the source could not unambiguously be matched to our data. The data are from *marinetraffic.com*. Deciles shown by row refer to the number of anchored ships in the port. In general, larger ports have stationary ships located farther from the port.

II.13 Nautical maps for dredging dummy variable

We obtained access to nautical maps of ports around the world from *marinetraffic.com*, see <https://www.marinetraffic.com/en/online-services/single-services/nautical-charts>. These detailed nautical charts have been constructed based on information from hydrographic organizations of different countries. They provide pilotage information including depth of water at high spatial disaggregation. Dredged channels are demarcated on these maps by a ‘safety contour’ that distinguishes the channel from the surrounding shallow waters (defined as less than 5 meters). We constructed a binary variable, ‘*Dredging*’, that takes the value 1 if a dredged channel is visible on the nautical chart in the 3-5 km buffer ring around the port.

II.14 Port cost data based on Blonigen and Wilson (2008)

Blonigen and Wilson (2008) estimate port costs as exporter-port fixed effects in a regression of bilateral HS 6-digit product level import charges that control for distance, value, value-to-weight, percentage of containerized traffic between the two ports, trade imbalances, time, product and importer-port fixed effects using U.S. census data for 1991 (see Blonigen and Wilson (2008) for additional details). The exporter fixed effects are all estimated relative to the port costs at Rotterdam. For our purposes, these relative measures need to be scaled to levels. We do this by setting the iceberg trade cost of passing through Rotterdam to be 1.004. This is based on estimates of revenue from handling one container to be approximately \$140 AUD (Australian Competition and Consumer Commission, 2017, p. 8) and the average value of a container to be 20,000 EUR (Kirchner, 2006, p. 4).¹⁴

Table S.4: Prediction of port cost

(1)	
Independent Variables	Port cost
ln(Shipment)	-0.033** (0.015)
Constant	0.444*** (0.145)
Observations	72
R-squared	0.074

Notes: The dependent variable, port cost, is taken from the port efficiency estimation in Blonigen and Wilson (2008) for 1991, available for 72 international port cities in our data (for details, see Appendix II.14). The regressor, $\ln(\text{Shipment})$, refers to our shipping data in 1990. Observations are weighted by the inverse of the squared standard error of the estimated port cost as given by Blonigen and Wilson (2008). Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

¹⁴These are industry-level averages as of 2016 (for revenue from container handling) and 2006 (for average value of cargo), and do not refer specifically to data from Rotterdam.

Table S.5: Predicted port cost

	N	mean	standard deviation
1950	2145	0.345	0.105
1960	2145	0.328	0.106
1970	2145	0.324	0.108
1980	2145	0.317	0.105
1990	2145	0.294	0.104

Notes: This table shows summary statistics for the predicted port cost based on the estimation in Table S.4. The 2,145 ports include the 553 ports with population data (Geopolis ports) as well as all other ports from the Lloyd's List data that do not have population data.

II.15 Data on frost-free days

We use data from the FAO GAEZ database (<http://www.fao.org/nr/gaez/en/>) to measure the average the number of frost-free days per year in each city. This database provides the average of this variable during the years between 1961 and 1990 in every cell over a 5 by 5 arc minute grid of the Earth. Using the geocoordinates of each city, we determine the grid cell in which the city is located, and assign the average number of frost-free days in the cell to the city.

II.16 Annual reports for ports

We were able to acquire annual reports for a number of port authorities in the United States during our sample period, 1950 to 1990, and for a handful of ports worldwide. Some ports have made historical annual reports available online, while for others, we have obtained the reports by contacting the port authorities. We use these reports for i) historical evidence (Section 1), ii) in the case of the Port of Seattle, to measure changes in land per unit of throughput during the period in which they containerized (Section 1), iii) to examine reporting on pollution and disamenities (Section 7), and iv) to calculate profit rates (Section I.7).

As accounting and reporting standards changed across ports and over time, we only kept ports that reported consistent information on profits over time (defined as revenue minus operating expenses and depreciation). These ports are: Houston, Los Angeles, Long Beach, New York/New Jersey, New Orleans, Seattle and Townsville (Australia). We tried to collect at least one observation per port for each decade between 1950 and 1990, and ended up with on average three decadal observations per port. The average

profit margin across all observations in our sample is 28%, with no clear time trend. Data sources are as follows;

[Houston](#). Port of Houston Authority of Harris County, Texas: ‘Comprehensive Annual Financial Report’ (various years). Thank you to Dollores Villareal at the Port of Houston for responding to our request and digitizing the data for us.

[Los Angeles](#). Port of Los Angeles Board of Harbor Commissioners: ‘Annual Report’ (various years). Thank you to Kurt Arendt at the Port of Los Angeles for responding to our request and sharing data.

[Long Beach](#). The Port of Long Beach California: ‘Harbor Highlights’ (various years). Accessible at <https://www.polb.com/port-info/history#historical-publications>.

[New York/New Jersey](#). The Port Authority of New York and New Jersey: ‘Annual Report’ (various years). These can be accessed online at <https://corpinfo.panynj.gov/pages/annual-reports/>.

[New Orleans](#). Board of Commissioners of the Port of New Orleans: ‘Annual Report Fiscal’ (various years). Thank you to Mandi Venderame at the Port of New Orleans for responding to our request and sharing data.

[San Francisco](#). The Port of San Francisco: ‘Annual Report’, other reports and planning maps from various years. Thank you to Randolph Quezada at the Port of San Francisco for numerous helpful conversations and for sharing scans.

[Seattle](#). The Port of Seattle: ‘Annual Report’ (various years) and planning maps. Thank you to Midori Okazaki, archivist at Puget Sound Regional Archives, for scanning the files during the COVID-19 lockdown while the archives were closed to the public.

[Townsville \(Australia\)](#). Townsville Harbor Board: ‘Report’ (various years). Thank you to the Port Authority for responding to our data request.

II.17 Maritime Silkroad Counterfactual, regression with country fixed effects

Table S.6 replicates Table B.16 of the online appendix, but adds country fixed effects.

References

- Abe, K. and J. Wilson (2009). *Weathering the Storm: Investing in Port Infrastructure to Lower Trade Costs in East Asia*. World Bank.
- Allen, T. and C. Arkolakis (2014). Trade and the Topography of the Spatial Economy. *Quarterly Journal of Economics* 129(3), 1085–1140.
- Allen, T. and C. Arkolakis (2019). The Welfare Effects of Transportation Infrastructure Improvements.

Table S.6: Maritime Silkroad Counterfactual, country fixed effects

	Baseline				Benchmark 1			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$\Delta \ln(\text{Shipment})$	$\Delta \ln(\text{Port cost})$	$\Delta \ln(\text{Market access})$	$\Delta \ln(\text{Population})$	$\Delta \ln(\text{Shipment})$	$\Delta \ln(\text{Port cost})$	$\Delta \ln(\text{Market access})$	$\Delta \ln(\text{Population})$
Treated port city	1.09629*** (0.10472)	-0.13053*** (0.01255)	0.01604*** (0.00277)	-0.02003** (0.00807)	0.91013*** (0.06377)	-0.10536*** (0.00000)	0.01290*** (0.00232)	0.01584*** (0.00285)
Untreated port city in treated country	-0.92756*** (0.03863)	0.00933 (0.00859)	0.04688*** (0.00595)	0.00551 (0.00953)	-0.86467*** (0.07237)	0.00000 (0.00000)	0.05045*** (0.00631)	-0.01007 (0.00775)
Port city in untreated country	0.00362*** (0.00045)	-0.00006* (0.00003)	-0.00140*** (0.00031)	-0.00133 (0.00108)	0.00262*** (0.00020)	0.00000 (0.00000)	-0.00134*** (0.00005)	0.00032*** (0.00007)
Inland city in treated country			0.05329*** (0.00482)	-0.00020 (0.00754)			0.05504*** (0.00516)	-0.00443 (0.00633)
Inland city in untreated country			-0.00192*** (0.00005)	0.00009 (0.00006)			-0.00155*** (0.00005)	0.00006 (0.00006)
Observations	553	544	2,636	2,636	553	553	2,636	2,636
R-squared	0.63923	0.38001	0.94168	0.17784	0.92353	1.00000	0.97896	0.28222

Notes: Treated port indicates the 24 treated ports of the Maritime Silkroad counterfactual. Treated country are countries that have at least one treated port. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Australian Competition and Consumer Commission (2017). Container Stevedoring and Monitoring Report 2016-17. Technical report, ACCC.

Baragwanath, K., R. Goldblatt, G. Hanson, and A. Khandelwal (2019). Detecting Urban Markets with Satellite Imagery: An Application to India. *Journal of Urban Economics* (103173).

Bernard, A., J. Eaton, J. Jensen, and S. Kortum (2003). Plants and Productivity in International Trade. *American Economic Review* 93(4), 1268–1290.

Blonigen, B. and W. Wilson (2008). Port Efficiency and Trade Flows. *Review of International Economics* 16(1), 21–36.

Ciccone, A. and R. Hall (1993). Productivity and the Density of Economic Activity. *National Bureau of Economic Research* (Working Paper 4313).

Desmet, K. and J. Rappaport (2017). The Settlement of the United States, 1800–2000: The Long Transition Towards Gibrat’s Law. *Journal of Urban Economics* 98, 50–68.

Donaldson, D. and A. Storeygard (2016). The View from Above: Applications of Satellite Data in Economics. *Journal of Economic Perspectives* 30(4), 171–98.

Haworth, R. B. (2020). Miramar ship index.

Kennan, J. and J. Walker (2011). The Effect of Expected Income on Individual Migration Decisions. *Econometrica* 79(1), 211–251.

Kirchner, M. (2006). Container Vessels and Risk Aggregation: The Cargo Underwriter’s View. Technical report, AXA Corporate Solutions.

Lloyd’s of London Press (1990). Lloyd’s Ports of the World.

Martín-Antón, M., d. C. J. M. Negro, V., J. S. López-Gutiérrez, and M. D. Esteban (2016). Review of coastal land reclamation situation in the world. *Journal of Coastal Research* 75, 667 – 671.

OECD (2018). The Belt and Road Initiative in the Global Trade, Investment, and Fi-

- nance Landscape. Technical report, OECD Publishing Paris.
- Port Authority of Seattle (1989). Seaport Properties Book. Technical report, Port Authority of Seattle.
- Port of Long Beach (1971). Harbor Highlights 1971 Report. Technical report, Port Authority of Long Beach.
- Port of New Orleans (1984). 88th Annual Report. Technical report, Port of New Orleans Board of Commissioners.
- Riffenburgh, R. (2012). A Project History of the Port of Long Beach 1970 to 2010. Technical report, Port of Long Beach.
- Saiz, A. (2010). The Geographic Determinants of Housing Supply. *Quarterly Journal of Economics* 125(3), 1253–1296.